

Garment Recovery with Shape and Deformation Priors

Ren Li Corentin Dumery Benoît Guillard Pascal Fua
 CVLab, EPFL

{name.surname}@epfl.ch

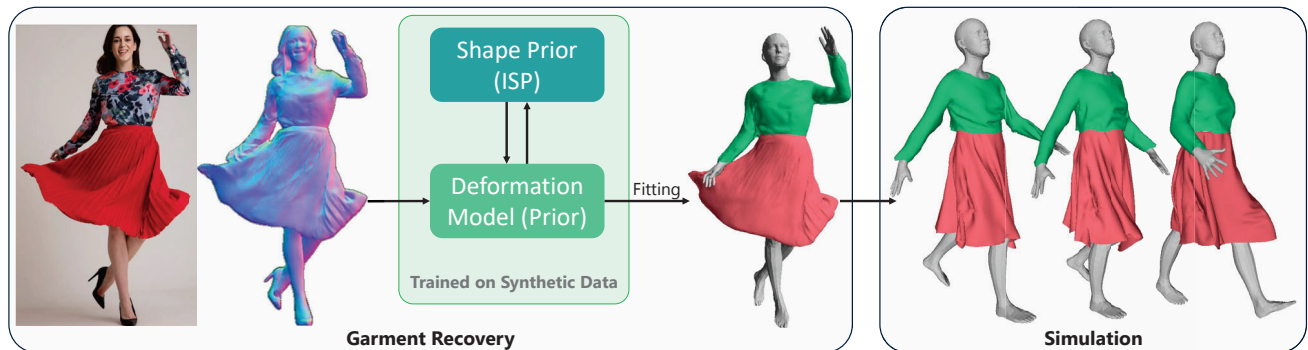


Figure 1. We propose a fitting method that leverages shape and deformation priors trained on synthetic data to recover realistic 3D garment mesh from in-the-wild images. We produce triangulated meshes that are directly usable for animation and simulation.

Abstract

While modeling people wearing tight-fitting clothing has made great strides in recent years, loose-fitting clothing remains a challenge. We propose a method that delivers realistic garment models from real-world images, regardless of garment shape or deformation. To this end, we introduce a fitting approach that utilizes shape and deformation priors learned from synthetic data to accurately capture garment shapes and deformations, including large ones. Not only does our approach recover the garment geometry accurately, it also yields models that can be directly used by downstream applications such as animation and simulation.

1. Introduction

3D clothed human recovery aims to reconstruct the body shape, pose, and clothing of many different people from images. It is key to applications such as fashion design, virtual try-on, 3D avatars, along with virtual and augmented reality. Recent years have seen tremendous progress in modeling people wearing tight-fitting clothing both in terms of body poses [6, 14, 21, 22, 25, 26, 28, 30, 38, 41, 54] and 3D shape of the clothes [5, 9–11, 20, 31, 32, 39]. However, loose-fitting clothing remains a challenge. Existing approaches either rely on mesh templates with limited generality, or produce models expressed in terms of 3D point clouds [37] or as a single watertight mesh that tightly binds the body

and garment together [53], neither of which is straightforward to integrate into downstream applications.

In this paper, we propose a method that overcomes these limitations and can effectively recover the shape of loose fitting garments from single images. The recovered garments can then be animated without any additional processing, as shown in Fig. 1. Starting from the Implicit Sewing Patterns (ISP) model [32] that represents garments in terms of a set of individual 2D panels and 3D surfaces associated to these panels, we introduce a deformation model that we apply to the 3D surfaces so that they can deviate substantially from the body shape. These deformations are conditioned on normals estimated from an input image of the target garment. They are learned from synthetic mesh data featuring loose clothing, where the deformations are taken to be those required to fit individual ISP 3D surfaces to the ground-truth 3D meshes.

Given the trained deformation model, we designed a two-stage fitting process to recover the 3D garment from in-the-wild images. First, the parameters of the pre-trained deformation model are optimized to produce a shape that minimizes the distance between garment outlines and segmented garment regions, the differences between garment normals and the normals estimated in the images by off-the-shelf-algorithms [46, 53], and a physics-based loss to promote physical plausibility of the results. Then, fine local details are recovered by directly optimizing the vertex positions of the reconstructed mesh with the same loss. Our fitting process does not require external 3D annotations, other

than the estimated normals of the target garment.

We demonstrate that our method can recover garments that go from tight- to loose-fitting and outperforms existing approaches [9, 11, 20, 32, 39, 53] in terms of reconstruction accuracy. Furthermore, our reconstructed meshes are directly usable for virtual try-on or animation, unlike [37, 53]. Our implementation and model weights are available at <https://github.com/liren2515/GarmentRecovery>.

2. Related Work

Before the advent of Deep Learning, garment shape recovery from images depended mostly on user defined outlines and shape-from-shading techniques [59]. Since then, data-driven techniques have become dominant.

Tight-Fitting Clothing. The majority of methods developed in recent years focus on clothing that clings relatively closely to the body. These can be classified into two main categories.

In the first category, are methods that model garments as surfaces that are distinct from the body surface and interact with it. DeepGarment [10], MGN [5], and BCNet [20] train neural networks on synthetic images to predict the vertex positions of specific mesh templates. [8] and [34] leverage normal estimation to optimize the vertex position of the template and recover wrinkle details. Such methods are inherently limited in the range of shapes they can handle, and those being trained on synthetic data can easily fail when facing real images. To overcome these limitations, SMPLicit [9], DIG [31], and ClothWild [39] leverage Signed Distance Functions (SDF) to recover a wide array of garment meshes from RGB images and the corresponding segmentation masks. However, to represent non-watertight garment surfaces using an SDF, one has to wrap around them a watertight surface with a minimum thickness, which reduces accuracy. This can be addressed by using Unsigned Distance Functions (UDFs) instead [11, 15] but creates robustness issues: if the UDF is even slightly inaccurate, the value of the surface is never exactly zero and holes can appear in the reconstructed models. In our experience, the Implicit Sewing Patterns (ISP) model of [32] effectively addresses the issues of generality, accuracy, and robustness. The garments consist of flat 2D panels whose boundary is defined by a 2D SDF. To each panel is associated a 3D surface parameterized by the 2D panel coordinates. Hence, different articles of clothing are represented in a standardized way, which allows the recovery of various garments from single images. This is why we choose it as the basis for our approach.

In the second category, are the many methods that represent body and garment using a single model. For example, in [19, 57] a volumetric regression network yields a voxel

representation of 3D clothed humans given a single image. Other works [3, 16, 45, 46] employ a pixel-aligned implicit function that defines 3D occupancy fields or signed distance fields for clothed humans. In [1, 2], displacement vectors or UV maps are used to represent deviations from a SMPL parametric body model [36]. Similarly, in [17, 18, 52, 58] parametric body models are combined with implicit representations to achieve robustness to pose changes. While effective, all these methods suffer from significant limitations, because they cannot separate the surface of the garment from that of the body, and they are at a disadvantage when it comes to modeling loose garments whose motion can be relatively independent from the body.

Loose-Fitting Clothing. There is a more limited number of methods designed to handle free-flowing garments. Some recent works [55, 60, 61] rely on complex physics simulation steps or feature line estimation to align the surface reconstruction with the input image. However, their dependence on garment templates limits their generality, in the same way it did for other template-based methods discussed above. Point-based methods that can reconstruct generic clothes have been proposed [50, 56] to overcome this. Unfortunately, point clouds are not straightforward to integrate into downstream applications. As a result, the method of [50] resorts to modified Poisson Surface Reconstruction (PSR) to create a garment surface from the point cloud, which can result in incorrect geometry. Another point-based representation is introduced in [37]. While being successful at modeling and animating humans wearing loose garments, it also relies on PSR to infer the mesh from the point cloud, yielding a single mesh that represents body and garment jointly. Furthermore, [37] does not explore how this representation can be fitted to images.

ECON [53] is a method specifically designed for clothed human recovery from images. By leveraging techniques such as normal integration and shape completion, it achieves visually appealing results for individuals wearing loose clothing. However, as [37], ECON produces a single watertight mesh that tightly binds the body and garment together, precluding easy use for applications such as cloth simulation and re-animation.

3. Method

Given an image of a clothed person and a body model extracted from it using existing techniques, our goal is to recover accurate 3D models of the garments matching the image. To this end, we add to the Implicit Sewing Pattern (ISP) garment model [32], which provides us with a shape prior for garment in its rest state, a deformation model that allows us to recover its potentially large deformations, as illustrated by Fig. 2. As a result, whereas the original ISP, like most current clothes-recovery algorithms, is limited to

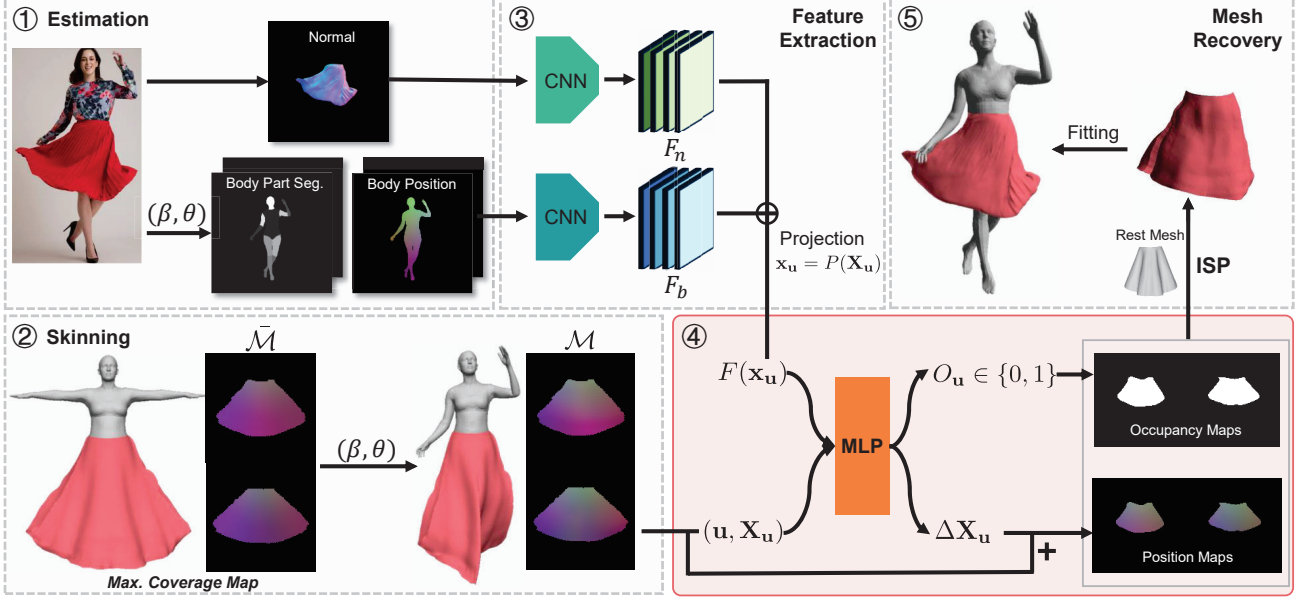


Figure 2. **Framework.** Given an image, (1) we first estimate the normal map of the target garment and the SMPL body parameters (β, θ) , which are used to compute the body part segmentation and position maps. (2) The maximum coverage garment shape $\bar{\mathcal{M}}$ is then skinned to closely fit to the body, yielding \mathcal{M} . Leveraging (3) pixel-aligned image features, our deformation model (4) predicts occupancy and position maps to correct \mathcal{M} for large deformations. (5) The 3D garment mesh is recovered using ISP and further refined.

tight-fitting clothing, our approach can handle both tight- and loose-fitting garments, such as skirts and open jackets.

In this section, we first describe briefly the ISP model upon which we build our approach. We then introduce the deformation model that underpins the main contribution of this paper, going from tight fitting clothes to loosely fitting and free flowing ones. Finally, we present our approach to fitting this model to real-world images.

3.1. ISP Garment Model

ISP is a garment model inspired by the sewing patterns that fashion designers use to represent clothes. A sewing pattern is made of several 2D panels along with information about how to stitch them into a complete garment. ISP implicitly models patterns using a 2D signed distance field and a 2D label field.

Formalization. Given the latent code \mathbf{z} of a garment and a point \mathbf{u} in the 2D UV space $\Omega = [-1, 1]^2$ of a 2D panel of that garment, ISP outputs the signed distance s to the panel boundary and a label l as

$$(s, l) = \mathcal{I}_{\Theta}(\mathbf{u}, \mathbf{z}), \quad (1)$$

where \mathcal{I}_{Θ} is a fully connected network. The zero crossing of the SDF defines the shape of the panel, with $s < 0$ indicating that \mathbf{u} is within the panel and $s > 0$ indicating that \mathbf{u} is outside the panel. The stitch information is encoded in l , where panel boundaries with the same label should be stitched together. To transform the 2D sewing patterns into

3D surfaces, a UV parameterization function \mathcal{A}_{Φ} is learned to perform the 2D-to-3D mapping

$$\mathbf{X} = \mathcal{A}_{\Phi}(\mathbf{u}, \mathbf{z}), \quad (2)$$

where $\mathbf{X} \in \mathbb{R}^3$ represents the 3D position of \mathbf{u} . Essentially, ISP registers each garment onto a unified 2D space Ω , and represents it using UV maps that record 2D-to-3D mapping, as shown in Fig. 3(b). Given the paired 2D sewing patterns and their 3D meshes, the pattern parameterization network \mathcal{I}_{Θ} and the UV parameterization network \mathcal{A}_{Φ} are trained by minimizing the losses

$$L_{\mathcal{I}} = L_{SDF} + L_{CE} + \|\mathbf{z}\|_2^2, \quad (3)$$

$$L_{\mathcal{A}} = L_{MSE} + L_{consist}, \quad (4)$$

respectively. L_{SDF} is the mean absolute error for the predicted SDF value s , L_{CE} is the cross-entropy loss for the predicted label l , L_{MSE} is the mean squared error of the predicted 3D position \mathbf{X} , and $L_{consist}$ is the loss to reduce the gap between the front and back panels. More details can be found in [32].

Training. Training ISP requires the 2D sewing patterns of 3D garments in a rest state, that is, draped on a T-pose neutral body. However, patterns are absent in most large scale garment datasets, including CLOTH3D [4]. We use the garment flattening algorithm of [42] to generate the required 2D patterns. The garment mesh is first divided into several pieces given predefined cutting rules. These pieces




Figure 3. **Cutting and flattening.** (a) The front (top) and back (bottom) surfaces after cutting. (b) The flattened panels and UV maps generated by ISP for (a). (c) The maximum-coverage UV maps and its represented 3D shape.

are then unfolded into 2D panels by minimizing an as-rigid-as-possible [33] energy, ensuring local area preservation between the 3D and 2D parameterizations. While doing this, we constrain the boundary vertices at places such as waist and sleeves to have a constant value along a specific axis. This enhances pattern consistency across the dataset. Fig. 3 illustrates this and shows the panels generated for a shirt. We provide more details in the supplementary material.

We generate a front and a back panels as the sewing pattern for each garment in our dataset (shirt, skirt, and trousers). Once ISP has been trained on these, we compute the maximum-coverage UV maps $\bar{\mathcal{M}}$ over the UV maps of each garment for each category, as shown in Fig. 3(c). The values of $\bar{\mathcal{M}}$ are taken to be

$$\bar{\mathcal{M}}[u, v] = \frac{\sum_i \mathbb{1}_{s_{\mathbf{u}}^i \leq 0} \cdot m_i[u, v]}{\sum_i \mathbb{1}_{s_{\mathbf{u}}^i \leq 0}}, \quad (5)$$

where m_i is the UV maps of garment i , $\mathbb{1}$ is the indicator function, $[\cdot, \cdot]$ denotes array addressing, and $s_{\mathbf{u}}^i$ is the SDF value of ISP at $\mathbf{u} = (u, v)$. The *maximum coverage map* $\bar{\mathcal{M}}$ encompasses information from all the patterns in the dataset. It represents the smallest possible map that covers all garments in a category, with the 3D position of each uv -point \mathbf{u} being the average of all garments that include \mathbf{u} . We use it as a prototype for a garment category and to compute an initial guess of the deformed garment given the body pose, as discussed below.

Skinning. As in many prior work [11, 31, 32], we use SMPL [36] to parameterize the body in terms of shape and pose parameters (β, θ) , and its extended skinning procedure for the 3D volume around the body to initially deform the 3D shape represented as $\bar{\mathcal{M}}$. More specifically, given a 2D point $\mathbf{u} = (u, v)$ in UV space Ω , we get its actual 3D position as $\bar{\mathbf{X}}_{\mathbf{u}} = \bar{\mathcal{M}}[u, v]$. We then deform it by computing

$$\begin{aligned} \mathbf{X}_{\mathbf{u}} &= W(\mathbf{X}_{(\beta, \theta)}, \beta, \theta, w(\bar{\mathbf{X}}_{\mathbf{u}})W), \\ \mathbf{X}_{(\beta, \theta)} &= \bar{\mathbf{X}}_{\mathbf{u}} + w(\bar{\mathbf{X}}_{\mathbf{u}})B_{\beta} + w(\bar{\mathbf{X}}_{\mathbf{u}})B_{\theta}, \end{aligned} \quad (6)$$

where $W(\cdot)$ is the SMPL skinning function with skinning weights $\mathcal{W} \in \mathbb{R}^{N_B \times 24}$, with N_B being the number of vertices of the SMPL body mesh, and $B_{\beta} \in \mathbb{R}^{N_B \times 3}$ and $B_{\theta} \in \mathbb{R}^{N_B \times 3}$ are the shape and pose displacements of SMPL, respectively. The diffused weights $w(\cdot) \in \mathbb{R}^{N_B}$ are computed by a neural network, which generalizes the SMPL skinning to any point in 3D space. By repeating this for all points in a panel, we obtain the vertex position map $\mathcal{M}[u, v] = \mathbf{X}_{\mathbf{u}}$.

3.2. Modeling Large Deformations

As shown in Fig. 2, the ISP approach described above generates a garment prototype \mathcal{M} that closely fits the underlying body. To a point $\mathbf{u} = (u, v)$ that belongs to a given garment panel, it associates the vertex position $\mathbf{X}_{\mathbf{u}} = \mathcal{M}[u, v]$ of Eq. 6, which is usually relatively close to the body. To model loose clothing, we now need to compute a potentially larger displacement $\Delta\mathbf{X}_{\mathbf{u}}$ to be added to $\mathbf{X}_{\mathbf{u}}$.

To evaluate $\Delta\mathbf{X}_{\mathbf{u}}$, we train a network \mathcal{D} that consists of an MLP to estimate the occupancy value and the corrective displacement value, and two CNNs to extract image features F_n and F_b from the image of normals and the segmentation and vertex position images of the SMPL body, respectively. Fig. 2 depicts this architecture. We obtain the pixel-aligned image feature for \mathbf{u} by computing

$$\begin{aligned} F(\mathbf{x}_{\mathbf{u}}) &= F_n \oplus F_b(\mathbf{x}_{\mathbf{u}}), \\ \mathbf{x}_{\mathbf{u}} &= P(\mathbf{X}_{\mathbf{u}}), \end{aligned} \quad (7)$$

where $P(\cdot)$ denotes the projection into the image and \oplus concatenation. The MLP takes as input \mathbf{u} , its 3D position and image features to predict the occupancy $O_{\mathbf{u}} \in \{0, 1\}$ and the corrective displacement $\Delta\mathbf{X}_{\mathbf{u}}$. By assembling the results of each point in the UV space, we obtain the final occupancy maps \mathcal{O} and the vertex position maps $\hat{\mathcal{M}}$, where $\hat{\mathcal{M}}[u, v] = \mathbf{X}_{\mathbf{u}} + \Delta\mathbf{X}_{\mathbf{u}}$. In essence, \mathcal{O} is the binarized SDF of the garment 2D panels, which implicitly defines the garment shape and geometry in the rest state, while $\hat{\mathcal{M}}$ encodes the deformed state for that garment. The 3D garment mesh

can be recovered from these 2D maps with ISP, as discussed below.

We use a single network \mathcal{D} for the front and the back panels, and encode the points \mathbf{u} of the front as $[u, v, +1]$ and those of the back as $[u, v, -1]$. We learn the parameters of \mathcal{D} by minimizing

$$L = \sum_{\mathbf{u} \in \Omega} \|\mathbf{X}_{\mathbf{u}} + \Delta \mathbf{X}_{\mathbf{u}} - \tilde{\mathbf{X}}_{\mathbf{u}}\|_2 + \lambda \sum_{\mathbf{u} \in \Omega} BE(O_{\mathbf{u}}, \tilde{O}_{\mathbf{u}}), \quad (8)$$

where $\tilde{\mathbf{X}}_{\mathbf{u}}$ and $\tilde{O}_{\mathbf{u}}$ are the ground truth vertex positions and occupancy values, $BE(\cdot)$ is the binary cross entropy, and λ is a weighting constant.

From 2D maps to 3D mesh. The occupancy value $O_{\mathbf{u}}$ indicates whether \mathbf{u} falls within the panels of the garment. To convert the generated 2D occupancy maps \mathcal{O} into a 3D garment mesh in its rest shape, we need to recover an ISP latent code \mathbf{z} as defined in Eq. 1. To this end, we find the vector \mathbf{z} such that the corresponding SDF of ISP best matches the produced occupancy maps by minimizing

$$\mathbf{z}^* = \operatorname{argmin}_{\mathbf{z}} \sum_{\mathbf{u} \in \Omega_-} R(s_{\mathbf{u}}(\mathbf{z})) + \sum_{\mathbf{u} \in \Omega_+} R(-s_{\mathbf{u}}(\mathbf{z})) + \lambda_{\mathbf{z}} \|\mathbf{z}\|_2, \quad (9)$$

where $\Omega_- = \{\mathbf{u} | O_{\mathbf{u}} = 1, \mathbf{u} \in \Omega\}$, $\Omega_+ = \{\mathbf{u} | O_{\mathbf{u}} = 0, \mathbf{u} \in \Omega\}$, $R(\cdot)$ is the ReLU function, $s_{\mathbf{u}}(\mathbf{z})$ is the SDF value of \mathbf{u} computed by ISP, and $\lambda_{\mathbf{z}}$ is the weighting constant. With \mathbf{z}^* , the rest garment mesh is inferred through ISP’s meshing and sewing process. The deformed garment mesh can be obtained by simply replacing the vertex position of the recovered mesh with the values stored in \mathcal{M} . This yields 3D garment meshes in both the rest and deformed states as shown in the top-right of Fig. 2, which are required for the application such as cloth simulation and can be used for further refinement as discussed below.

3.3. Fitting the Models to Images

For practical reasons, the range of garment materials, external forces, and body motions present in the training data is limited. As a result, given in-the-wild images as input, the trained model can produce inaccurate results as shown in Fig. 4(b). To remedy this and to leverage the deformation prior that the network \mathcal{D} captures, we refine the result by minimizing a loss function with respect to the pretrained deformation parameters of \mathcal{D} as in [13, 21, 51]. L is designed to promote a good match between the garment mesh and image observations. We take it to be

$$L = \lambda_C L_{CD} + \lambda_n L_{normal} + \lambda_p L_{physics}, \quad (10)$$

$$L_{CD} = d(x_c^f, x_{I_n}), \quad (11)$$

$$L_{normal} = \sum_{i \in f} 1 - \cos(n^i, I_n(x_c^i)), \quad (12)$$

$$L_{physics} = L_{strain} + L_{bend} + L_{gravity} + L_{col}, \quad (13)$$




Figure 4. **Fitting results.** Given (a) the normal estimation of an in-the-wild image, (b) is the inference result with ISP recovered geometry. (c) is obtained by optimizing the parameters of the pretrained deformation model. Further refinement of the mesh vertex positions yields (d).

where x_c^f is the 2D projection of the centers of visible faces f after mesh rasterization, x_{I_n} denotes the coordinates of foreground pixels, n^i is the normal of face i , $I_n(x_c^i)$ is the normal image values at x_c^i , and λ_C , λ_n and λ_p are the balancing scalars. $d(\cdot)$ and $\cos(\cdot)$ are the functions measuring the 2D Chamfer Distance and the cosine similarity, respectively. $L_{physics}$ is a physics-based loss derived from [40, 48], which computes the membrane strain energy L_{strain} caused by the deformation, the bending energy L_{bend} resulting from the folding of adjacent faces, the gravitational potential energy $L_{gravity}$ and the penalty for body-garment collision L_{col} . Minimizing L_{CD} induces an external force of stretching or compression on the garment mesh to align its 2D projection with the given image, while minimizing $L_{physics}$ ensures that the mesh exhibits physically plausible deformation adhering to the shape constraints of the rest-state mesh recovered by ISP.

Minimizing the loss L with respect to the deformation parameters yields a mesh whose overall shape matches the input image, as illustrated in Fig. 4(c). However, since neural networks tend to learn low-frequency functions [43], the result might be too smooth. To recover fine surface details, we perform a refinement step by minimizing L directly with respect to the coordinates of the garment mesh vertices. This generates realistic local details, such as wrinkles on the surface, as shown in Fig. 4(d).

4. Experiments

In short, our method begins by inferring the shape and the deformation of the garment in terms of the occupancy and position maps. Leveraging the shape prior of ISP, we then

recover the garment geometry from the occupancy maps and deform it using the position maps. Next, we refine the initial deformed mesh to better align with image observations by fine-tuning the pre-trained network \mathcal{D} that captures the deformation prior. Finally, we recover fine details through vertex-level optimization of the garment mesh. In this section, we demonstrate the effectiveness of this process and compare it to that of other state-of-the-art methods.

4.1. Implementation Details

Following the implementation of [32], both the pattern parameterization model \mathcal{L}_Θ and the UV parameterization model \mathcal{A}_Φ of ISP have two separate MLPs for the front and back panels. Each MLP has 7 layers with Softplus activations. The dimension of latent code \mathbf{z} is 32. The skinning weight model w is a 9-layer MLP with leaky ReLU activations, whose output is normalized by a final Softmax layer. \mathcal{L}_Θ and \mathcal{A}_Φ are trained jointly for 9000 iterations with a batch size of 50. w is trained with the same parameters as [32]. For the image feature extraction, we use two separate ConvNeXt [35] networks to extract multi-scale garment and body features of sizes 96, 192, 384, 768, which are concatenated as the final features. The point UV coordinates \mathbf{u} , 3D position $\mathbf{X}_\mathbf{u}$, and image features $F(\mathbf{x}_\mathbf{u})$ are projected separately to 384 dimensions by three linear layers, which are then concatenated as the input of the MLP of the deformation model. The MLP of the deformation model has 10 layers with a skip connection from the input layer to the middle, and uses Gaussian functions as the activation layer following [44]. The CNNs and MLP of the deformation model are trained jointly for 40 epochs with the Adam optimizer [23] and a learning rate of 10^{-4} . For real images, we use [53] and [12] to obtain their normal and SMPL body parameter estimations, respectively. The garment segmentation masks are generated by leveraging the segmentation of SAM [24] and the semantic labels of [29].

4.2. Dataset, Evaluation Metrics, and Baseline

Our models are trained on CLOTH3D [4], which is a synthetic dataset with motion sequences of 3D clothed human. It contains garment in the rest state, and deformed states caused by the motion of underlying body. For each garment, it has a single simulated sequence up to 10 seconds. It covers a large variety of garment in different shapes, types and topologies. For the training of ISP, we randomly select 400 shirts, 200 skirts and 200 pairs of trousers, and generate their sewing patterns by the method described in Sec. 3.1. The deformation model is trained on the corresponding simulated sequences. For each frame, we render 11 normal images for the garment mesh with random rotations around the Y-axis, which produces 40K, 40K and 20K training images for shirt, skirt and trousers respectively. During training, we augment the data with image flipping and rotation.

To evaluate the garment reconstruction quality, we use the Chamfer Distance (CD) between the ground truth and the recovered garment mesh, and Intersection over Union (IoU) between the ground truth mask and the rendered mask of reconstructed garment mesh.

We compare our method against state-of-the-art methods BCNet [20], SMPLicit [9], ClothWild [39], DrapeNet [11] and ISP [32]. SMPLicit, DrapeNet and ISP use the garment segmentation mask for reconstruction, while BCNet and ClothWild take the RGB images as input.

4.3. Comparison with State-of-the-Art Methods

CD ($\times 10^3$) \downarrow	Skirt	Shirt	Trousers	IoU \uparrow	Skirt	Shirt	Trousers
SMPLicit	8.05	3.49	0.81	SMPLicit	0.657	0.527	0.782
DrapeNet	n/a	1.54	0.84	DrapeNet	n/a	0.715	0.845
ISP	7.26	2.02	0.91	ISP	0.709	0.665	0.781
Ours	2.87	0.44	0.28	Ours	0.940	0.939	0.946
<i>Ours-GT</i>	<i>1.67</i>	<i>0.31</i>	<i>0.23</i>	<i>Ours-GT</i>	<i>0.953</i>	<i>0.942</i>	<i>0.954</i>

Table 1. **Quantitative comparisons.** Our method outperforms SMPLicit, DrapeNet, and ISP in terms of CD and IoU on all three garment categories, as shown in the second-to-last row of both tables (Ours). These results were obtained using normals estimated from the images using [53]. In the last row (Ours-GT), we provide the results we obtained using ground-truth normals instead.

Quantitative Results. Due to the absence of publicly available real dataset for evaluating garment reconstruction, we utilize synthetic data consisting of 30 unseen shirts, 30 unseen skirts, and 30 pairs of unseen trousers from CLOTH3D for quantitative evaluation. We compare our method with SMPLicit, DrapeNet, and ISP, which recover garment meshes from segmentation masks. The segmentation masks required by the baselines are rendered from the ground truth mesh data, while the normal images required by our method are estimated with [53]. As shown in Tab. 1, our approach significantly outperforms the baselines in terms of Chamfer Distance (CD) and Intersection over Union (IoU) across all garment categories.

For comparison purposes, we re-ran our algorithm using the ground-truth normals and report the results in the last row of Tab. 1. As expected, this leads to an improvement in reconstruction accuracy, but the increase is only modest. This highlights the robustness of our approach to the slight inaccuracies that can be expected from a normal-estimation algorithm.

Qualitative Results. Fig. 5 shows the qualitative comparison for the results reconstructed from in-the-wild images. Since BCNet is trained only on synthetic RGB data, it is not able to predict accurate body and garment results. Both SMPLicit and ClothWild generate unrealistic watertight meshes that do not accurately represent real garments.

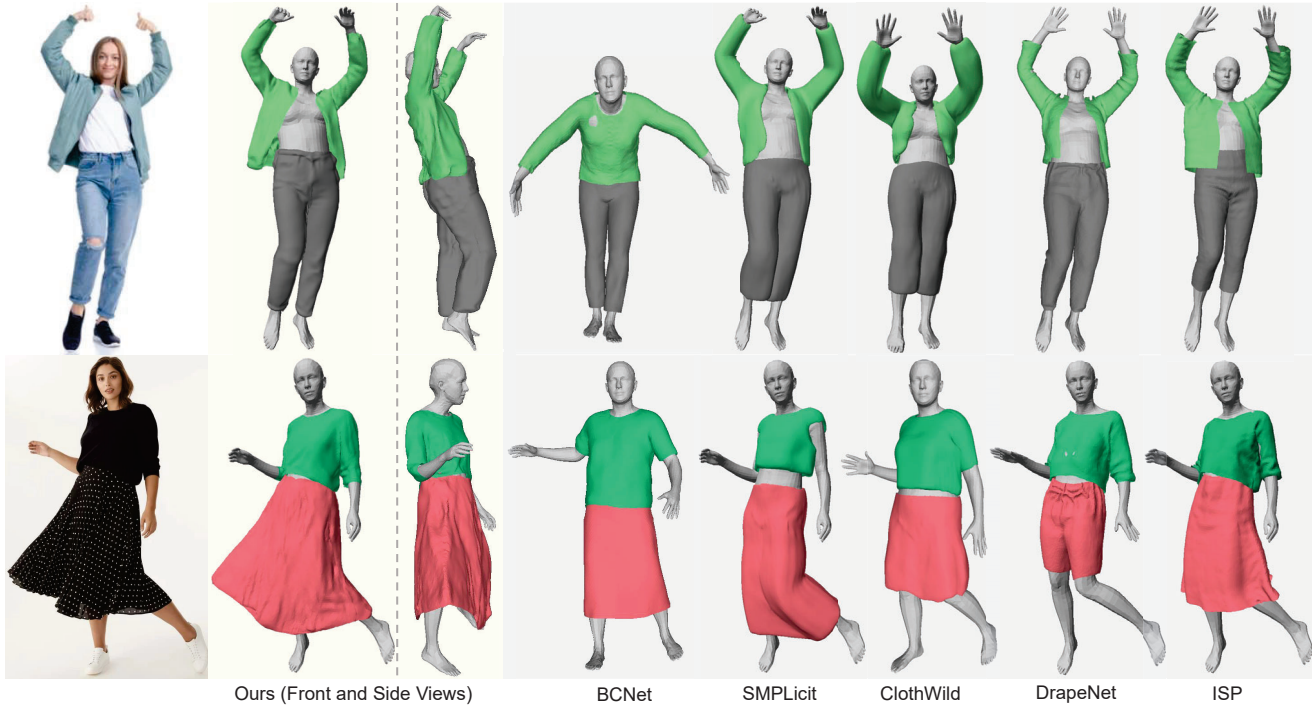


Figure 5. **Comparison against SOTA methods.** From left to right, we show the input image and the 3D garment meshes recovered by our method and SOTA methods: BCNet, SMPLicit, ClothWild, DrapeNet, ISP. (Since skirt is unavailable for DrapeNet, a random pair of trousers is put on its result in the second row.)

CD ($\times 10^3$) \downarrow	GT	Est.	IoU \uparrow	GT	Est.
$(-D^*, -v^*)$	2.04	3.58	$(-D^*, -v^*)$	0.848	0.751
$(-D^*, +v^*)$	1.80	3.39	$(-D^*, +v^*)$	0.947	0.904
$(+D^*, -v^*)$	1.75	2.93	$(+D^*, -v^*)$	0.941	0.923
$(+D^*, +v^*)$	1.67	2.87	$(+D^*, +v^*)$	0.953	0.940

Table 2. **Ablation study.** GT and Est. mean using ground-truth and estimated normal images as input, respectively. $+/-D^*$ denote with/without the finetuning of the deformation model. $+/-v^*$ denote with/without the optimization of vertex positions.

While DrapeNet and ISP are able to recover open surfaces for garments, their results remain closely adhered to the body, similar to SMPLicit and ClothWild. In contrast, our method can faithfully recover garment mesh from input images, as clearly observed in the depiction of the open jacket and the flowing long skirt that stand away from the body.

4.4. Ablation Study

Table 2 presents the evaluation results of our method with and without fitting on the test set of skirts. The results of row 1, 3 and 4 indicate that optimizing the parameters of the pretrained deformation network $(+D^*, -v^*)$ improves the quality of the raw inference results, and optimizing the vertex positions $(+D^*, +v^*)$ further enhances the outcome,




Figure 6. **Fitting strategy** comparison between (a) our fitting method and (b) only vertex optimization.

yielding the lowest CD and highest IoU.

However, as demonstrated in the second row, directly optimizing the vertex positions on the raw inference results without optimizing the deformation model $(-D^*, +v^*)$ is not as effective. The deformation model encapsulates a continuous deformation field. When its weights are optimized based on partial observations, the entire field undergoes modification, thereby influencing all mesh vertices globally. On the contrary, a direct vertex optimization with partial observations predominantly affects the mesh vertices locally. While this can capture localized details like wrinkles, it struggles to resolve discrepancies in the overall shape. Fig. 6 displays the results of our methods and an ablation with




Figure 7. **Recovering from in-the-wild images.** Our method is able to recover realistic meshes for garments with diverse shapes and deformations.

vertex optimization only. Unsurprisingly, the latter fails at recovering the correct shape and produces implausible deformation on the mesh surface.

4.5. More Results

Fig. 7 shows a collection of results reconstructed by our method from in-the-wild images. Our method can produce realistic 3D meshes with fine details across a wide range of garment types, from tight-fitting attire to more relaxed and flowing outfits.

5. Conclusion

We have presented a novel approach to recovering realistic 3D garment meshes from in-the-wild images featuring loose fitting clothing. It relies on a fitting process that imposes shape and deformation priors learned on synthetic data to accurately capture garment shape and deformations. In future work, we will extend our approach to modeling deformations over time from video sequences while enforcing temporal consistency of the reconstructions.

Acknowledgement. This project was supported in part by the Swiss National Science Foundation.

References

- [1] T. Alldieck, M. Magnor, B. Bhatnagar, C. Theobalt, and G. Pons-Moll. Learning to reconstruct people in clothing from a single RGB camera. In *Conference on Computer Vision and Pattern Recognition*, pages 1175–1186, 2019. 2
- [2] T. Alldieck, G. Pons-Moll, C. Theobalt, and M. Magnor. Tex2shape: Detailed full human body geometry from a single image. In *International Conference on Computer Vision*, pages 2293–2303, 2019. 2
- [3] T. Alldieck, M. Zanfir, and C. Sminchisescu. Photorealistic monocular 3d reconstruction of humans wearing clothing. In *Conference on Computer Vision and Pattern Recognition*, pages 1506–1515, 2022. 2
- [4] H. Bertiche, M. Madadi, and S. Escalera. CLOTH3D: Clothed 3D Humans. In *European Conference on Computer Vision*, pages 344–359, 2020. 3, 6, 1
- [5] B. L. Bhatnagar, G. Tiwari, C. Theobalt, and G. Pons-Moll. Multi-Garment Net: Learning to Dress 3D People from Images. In *International Conference on Computer Vision*, 2019. 1, 2
- [6] F. Bogo, A. Kanazawa, C. Lassner, P. Gehler, J. Romero, and M. J. Black. Keep It SMPL: Automatic Estimation of 3D Human Pose and Shape from a Single Image. In *European Conference on Computer Vision*, 2016. 1
- [7] R. Brouet, A. Sheffer, L. Boissieux, and M.P. Cani. Design Preserving Garment Transfer. *ACM Transactions on Graphics*, 31(4):361–3611, 2012. 1
- [8] A. Casado-Elvira, M. Trinidad, and D. Casas. PERGAMO: Personalized 3d garments from monocular video. In *Computer Graphics Forum*, pages 293–304, 2022. 2
- [9] E. Corona, A. Pumarola, G. Alenya, G. Pons-Moll, and F. Moreno-Noguer. Smplicit: Topology-Aware Generative Model for Clothed People. In *Conference on Computer Vision and Pattern Recognition*, 2021. 1, 2, 6
- [10] R. Danerek, E. Dibra, C. öztireli, R. Ziegler, and M. Gross. Deepgarment : 3D Garment Shape Estimation from a Single Image. *Eurographics*, 2017. 2
- [11] Luca DeLuigi, Ren Li, Benoît Guillard, Mathieu Salzmann, and Pascal Fua. DrapeNet: Generating Garments and Draping them with Self-Supervision. In *Conference on Computer Vision and Pattern Recognition*, 2023. 1, 2, 4, 6
- [12] Y. Feng, V. Choutas, T. Bolkart, D. Tzionas, and M. J. Black. Collaborative Regression of Expressive Bodies using Moderation. In *International Conference on Computer Vision*, 2021. 6
- [13] M. Gadelha, R. Wang, and S. Maji. Deep manifold prior. In *International Conference on Computer Vision*, pages 1107–1116, 2021. 5
- [14] G. Georgakis, R. Li and S. Karanam, T. Chen, J. Kořecká, and Z. Wu. Hierarchical kinematic human mesh recovery. In *European Conference on Computer Vision*, pages 768–784, 2020. 1
- [15] B. Guillard, F. Stella, and P. Fua. MeshUDF: Fast and Differentiable Meshing of Unsigned Distance Field Networks. In *European Conference on Computer Vision*, 2022. 2
- [16] T. He, J. Collomosse, H. Jin, and S. Soatto. Geo-pifu: Geometry and pixel aligned implicit functions for single-view human reconstruction. *Advances in Neural Information Processing Systems*, 33:9276–9287, 2020. 2
- [17] T. He, Y. Xu, S. Saito, S. Soatto, and T. Tung. Arch++: Animation-ready clothed human reconstruction revisited. In *International Conference on Computer Vision*, pages 11046–11056, 2021. 2
- [18] Z. Huang, Y. Xu, C. Lassner, H. Li, and T. Tung. Arch: Animatable Reconstruction of Clothed Humans. In *Conference on Computer Vision and Pattern Recognition*, pages 3093–3102, 2020. 2
- [19] A. Jackson, C. Manafas, and G. Tzimiropoulos. 3d human body reconstruction from a single image via volumetric regression. In *European Conference on Computer Vision Workshops*, 2018. 2
- [20] B. Jiang, J. Zhang, Y. Hong, J. Luo, L. Liu, and H. Bao. Bcnet: Learning body and cloth shape from a single image. In *European Conference on Computer Vision*, 2020. 1, 2, 6
- [21] H. Joo, N. Neverova, and A. Vedaldi. Exemplar Fine-Tuning for 3D Human Pose Fitting Towards In-the-Wild 3D Human Pose Estimation. In *International Conference on 3D Vision*, 2020. 1, 5
- [22] A. Kanazawa, M. J. Black, D. W. Jacobs, and J. Malik. End-To-End Recovery of Human Shape and Pose. In *Conference on Computer Vision and Pattern Recognition*, 2018. 1
- [23] D. P. Kingma and J. Ba. Adam: A Method for Stochastic Optimisation. In *International Conference on Learning Representations*, 2015. 6
- [24] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. Berg, W. Lo, P. Dollár, and R. Girshick. Segment anything. In *International Conference on Computer Vision*, 2023. 6
- [25] N. Kolotouros, G. Pavlakos, M. J. Black, and K. Daniilidis. Learning to Reconstruct 3D Human Pose and Shape via Model-Fitting in the Loop. In *Conference on Computer Vision and Pattern Recognition*, 2019. 1
- [26] C. Lassner, J. Romero, M. Kiefel, F. Bogo, M.J. Black, and P.V. Gehler. Unite the People: Closing the Loop Between 3D and 2D Human Representations. In *Conference on Computer Vision and Pattern Recognition*, 2017. 1
- [27] B. Lévy, S. Petitjean, N. Ray, and J. Maillot. Least Squares Conformal Maps for Automatic Texture Atlas Generation. In *ACM SIGGRAPH*, 2002. 1
- [28] J. Li, C. Xu, Z. Chen, S. Bian, L. Yang, and C. Lu. Hybrik: A hybrid analytical-neural inverse kinematics solution for 3d human pose and shape estimation. In *Conference on Computer Vision and Pattern Recognition*, 2021. 1
- [29] P. Li, Y. Xu, Y. Wei, and Y. Yang. Self-Correction for Human Parsing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(6):3260–3271, 2020. 6
- [30] R. Li, M. Zheng, S. Karanam, T. Chen, and Z. Wu. Everybody Is Unique: Towards Unbiased Human Mesh Recovery. In *British Machine Vision Conference*, 2021. 1
- [31] R. Li, B. Guillard, E. Remelli, and P. Fua. DIG: Draping Implicit Garment over the Human Body. In *Asian Conference on Computer Vision*, 2022. 1, 2, 4
- [32] Ren Li, Benoît Guillard, and Pascal Fua. ISP: Multi-Layered Garment Draping with Implicit Sewing Patterns. In *Ad-*

- vances in Neural Information Processing Systems*, 2023. 1, 2, 3, 4, 6
- [33] L. Liu, L. Zhang, Y. Xu, C. Gotsman, and S. J. Gortler. A Local/Global Approach to Mesh Parameterization. In *Proceedings of the Symposium on Geometry Processing*, pages 1495–1504, 2008. 4, 1
- [34] X. Liu, J. Li, and G. Lu. Modeling Realistic Clothing from a Single Image under Normal Guide. *IEEE Transactions on Visualization and Computer Graphics*, 2023. 2
- [35] Z. Liu, H. Mao, C. Wu, C. Feichtenhofer, T. Darrell, and S. Xie. A ConvNet for the 2020s. In *Conference on Computer Vision and Pattern Recognition*, 2022. 6
- [36] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M.J. Black. SMPL: A Skinned Multi-Person Linear Model. *ACM SIGGRAPH Asia*, 34(6), 2015. 2, 4
- [37] Qianli Ma, Jinlong Yang, Michael J. Black, and Siyu Tang. Neural Point-based Shape Modeling of Humans in Challenging Clothing. In *International Conference on 3D Vision*, 2022. 1, 2
- [38] G. Moon and K.M. Lee. I2L-MeshNet: Image-to-Lixel Prediction Network for Accurate 3D Human Pose and Mesh Estimation from a Single RGB Image. In *European Conference on Computer Vision*, 2020. 1
- [39] G. Moon, H. Nam, T. Shiratori, and K.M. Lee. 3d clothed human reconstruction in the wild. In *European Conference on Computer Vision*, 2022. 1, 2, 6
- [40] R. Narain, A. Samii, and J.F. O’Brien. Adaptive anisotropic remeshing for cloth simulation. *ACM Transactions on Graphics*, 2012. 5
- [41] M. Omran, C. Lassner, G. Pons-Moll, P. Gehler, and B. Schiele. Neural Body Fitting: Unifying Deep Learning and Model-Based Human Pose and Shape Estimation. In *International Conference on 3D Vision*, 2018. 1
- [42] N. Pietroni, C. Dumery, R. Falque, M. Liu, T. Vidal-Calleja, and O. Sorkine-Hornung. Computational pattern making from 3D garment models. *ACM Transactions on Graphics*, 41(4):1–14, 2022. 3, 1
- [43] N. Rahaman, A. Baratin, D. Arpit, F. Draxler, M. Lin, F. Hamprecht, Y. Bengio, and A. Courville. On the spectral bias of neural networks. In *International Conference on Machine Learning*, 2019. 5
- [44] S. Ramasinghe and S. Lucey. Beyond Periodicity: Towards a Unifying Framework for Activations in Coordinate-MLPs. In *European Conference on Computer Vision*, pages 142–158, 2022. 6
- [45] S. Saito, Z. Huang, R. Natsume, S. Morishima, A. Kanazawa, and H. Li. PIFu: Pixel-Aligned Implicit Function for High-Resolution Clothed Human Digitization. In *International Conference on Computer Vision*, 2019. 2
- [46] S. Saito, T. Simon, J. Saragih, and H. Joo. Pifuhd: Multi-Level Pixel-Aligned Implicit Function for High-Resolution 3D Human Digitization. In *Conference on Computer Vision and Pattern Recognition*, pages 84–93, 2020. 1, 2
- [47] I. Santesteban, N. Thuerey, M. A. Otaduy, and D. Casas. Self-Supervised Collision Handling via Generative 3D Garment Models for Virtual Try-On. In *Conference on Computer Vision and Pattern Recognition*, 2021. 1
- [48] I. Santesteban, M.A. Otaduy, and D. Casas. SNUG: Self-Supervised Neural Dynamic Garments. In *Conference on Computer Vision and Pattern Recognition*, 2022. 5
- [49] O. Sorkine-Hornung and M. Rabinovich. Least-squares rigid motion using svd. *Computing*, 1(1):1–5, 2017. 1
- [50] A. Srivastava, C. Pokhariya, S. Jinka, and A. Sharma. xCloth: Extracting Template-free Textured 3D Clothes from a Monocular Image. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 2504–2512, 2022. 2
- [51] D. Ulyanov, A. Vedaldi, and V. Lempitsky. Deep Image Prior. In *Conference on Computer Vision and Pattern Recognition*, pages 9446–9454, 2018. 5
- [52] Y. Xiu, J. Yang, D. Tzionas, and M. J Black. Icon: Implicit clothed humans obtained from normals. In *Conference on Computer Vision and Pattern Recognition*, pages 13286–13296, 2022. 2
- [53] Y. Xiu, J. Yang, X. Cao, D. Tzionas, and M. J Black. ECON: Explicit Clothed humans Optimized via Normal integration. In *Conference on Computer Vision and Pattern Recognition*, pages 512–523, 2023. 1, 2, 6, 3, 4
- [54] F. Yang, R. Li, G. Georgakis, S. Karanam, T. Chen, H. Ling, and Z. Wu. Robust multi-modal 3d patient body modeling. In *Conference on Medical Image Computing and Computer Assisted Intervention*, pages 86–95, 2020. 1
- [55] S. Yang, Z. Pan, T. Amert, K. Wang, L. Yu, T. Berg, and M. Lin. Physics-inspired garment recovery from a single-view image. *ACM Transactions on Graphics*, 37(5):1–14, 2018. 2
- [56] I. Zakharkin, K. Mazur, A. Grigorev, and V. Lempitsky. Point-based modeling of human clothing. In *International Conference on Computer Vision*, 2021. 2
- [57] Z. Zheng, T. Yu, Y. Wei, Q. Dai, and Y. Liu. Deephuman: 3d human reconstruction from a single image. In *International Conference on Computer Vision*, pages 7739–7749, 2019. 2
- [58] Z. Zheng, T. Yu, Y. Liu, and Q. Dai. Pamir: Parametric model-conditioned implicit representation for image-based human reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(6):3170–3184, 2021. 2
- [59] B. Zhou, X. Chen, Q. Fu, K. Guo, and P. Tan. Garment modeling from a single image. In *Computer Graphics Forum*, pages 85–91, 2013. 2
- [60] H. Zhu, Y. Cao, H. Jin, W. Chen, D. Du, Z. Wang, S. Cui, and X. Han. Deep fashion3d: A dataset and benchmark for 3d garment reconstruction from single images. In *European Conference on Computer Vision*, pages 512–530, 2020. 2
- [61] H. Zhu, L. Qiu, Y. Qiu, and X. Han. Registering explicit to implicit: Towards high-fidelity garment mesh reconstruction from single images. In *Conference on Computer Vision and Pattern Recognition*, pages 3845–3854, 2022. 2